

1                   In the Claims:

2                   1. (previously presented) A method for processing a database query,  
3                   comprising:

4                   partially pre-aggregating records in a database according to a single  
5                   grouping column to provide a result that contains at least two records having like  
6                   grouping column values;

7                   aggregating records derived from the partial pre-aggregation to provide a  
8                   result that contains records having unique grouping column values; and

9                   estimating the costs and benefits of the partial pre-aggregation, and partially  
10                  pre-aggregating the records only if the estimating indicates that the benefits are  
11                  greater than the costs.

12                  2. (Original.) The method as recited in claim 1, wherein the partially  
13                  pre-aggregating further comprises:

14                  maintaining a record store in memory, the record store having one record  
15                  for each different grouping column value encountered in the operation;

16                  receiving a new record;

17                  combining the new record with a record having the same grouping column  
18                  value, if such a record exists; and

19                  adding the new record to the record store in the memory if there is no record  
20                  in the record store that has the same grouping column value as the new record.

1           3. (Original.) The method as recited in claim 2, further comprising:  
2           adding additional new records to the record store until the record store  
3           reaches a capacity such that it can accept no new records; and  
4           outputting one or more records from the record store to a subsequent  
5           database operator.

6  
7           4. (Original.) The method as recited in claim 3, wherein after the one or  
8           more records have been output to the subsequent database operator, the adding and  
9           outputting are repeated until there are no new records to process.

10  
11          5. (Original.) The method as recited in claim 4, wherein any records  
12          remaining in the record store after there are no new records to process are output to  
13          the subsequent database operator.

14  
15          6. (Original.) The method as recited in claim 3, wherein the subsequent  
16          database operator is a join.

17  
18          7. (Cancelled.)

19  
20          8. (Original.) The method as recited in claim 1, wherein the partially  
21          pre-aggregating includes utilizing a hashing function.

22  
23  
24  
25

1       9. (Original.) The method as recited in claim 1, wherein the partial pre-  
2 aggregating creates a record store in memory, and wherein the method further  
3 comprises utilizing the record store in memory for one or more other database  
4 operators.

5       10. (Cancelled.)

6       11. (Original) A computer programmed to perform the method recited in  
7       claim 1.

8       12—23. (Cancelled.)

1       24. (currently amended) A relational database computer program stored  
2       on a computer-readable medium, the relational database computer program  
3       comprising computer-executable instructions that, when executed on a computer,  
4       perform the following steps:

5       receiving a stream of input records;

6       partially pre-aggregating the input records according to a single grouping  
7       column to provide a result that contains at least two records having like grouping  
8       column values;

9       aggregating the input records in the stream according to a single grouping  
10      column as it is received to create a record store;

11      joining records in the record store the partially pre-aggregated records with  
12      other data to create a record store; and

13      aggregating records within the record store to provide a result that contains  
14      records having unique grouping column values.

15      aggregating the records output from the join;

16      determining if it is optimal to aggregate the input records prior to  
17      performing the join;

18      performing the aggregation prior to the join only if a determination is made  
19      that it is optimal to perform an aggregation prior to the join; and

20      wherein the records output from the join include at least two records that  
21      have an identical grouping column value in the single grouping column.

1       25. (Original.) The relational database computer program as recited in  
2       claim 24, wherein:

3               the record store has a capacity that is less than the number of records in the  
4       stream of input records; and

5               the aggregating each input record is performed until the record store reaches  
6       capacity.

7       26. (Cancelled.)

8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25